

Collaboration through competition

Community-run bioinformatics challenges not only assess the state of the art but can help advance it.

Scientists have learned over the years that one of the best ways to bring people together is to have them compete. Competitions focus the community on a common problem and provide a rigorous evaluation of the most promising solutions. To accelerate research and spur innovation, organizers and researchers should continue to champion enthusiastic engagement as well as support efforts and newer formats that make competitions more interactive.

Bioinformatics challenges cover a wide range of subjects, from predicting protein structures to defining objects in noisy images. Teams are given access to unpublished benchmark data and typically have a few months to return their best predictions and protocols for expert evaluation and ranking.

The contests are not just exercises to crown a winner, like naming the best chef or top vocalist in a popular contest. They are also opportunities to define the goals of a field, clarify community standards and identify bottlenecks. To help find consensus, end users and developers should be included in the planning stages. For example, the choice of metrics is critical for judging how well a tool performs. In the Assemblathon competition for genome assembly, organizers established the primary metrics by an empirical process and then considered additional possibilities from anyone in the community, so long as he or she volunteered to analyze them.

Many of the benefits of competition are reaped at the end, when researchers gather to discuss triumphs and failures, learn from each other's methods and form collaborations. There is always a danger that the intense activity generated by the contest will dissipate after the event finishes. To keep the pressure on, one strategy is to iterate contests regularly. The Critical Assessment of Genome Interpretation (CAGI) community experiments on genomic variation interpretation are run annually on related themes, giving teams some urgency to develop new ideas before the next event.

There is also growing evidence that innovation can be extracted in a more immediate and direct way from the talent assembled on the contest's playing field. When the fifth Dialogue for Reverse Engineering Assessments and Methods (DREAM) challenge ended, the organizers merged leading tools and systematically tested these hybrid methods. The hybrid approaches performed

well overall and turned out to be more robust than individual methods across testing categories, confirming a body of work on ensemble methods in machine learning. The results argue for the need for some collaborative competition formats.

In this spirit, a number of newer DREAM challenges use an innovative two-step framework. After a traditional competition and evaluation, developers from a handful of the best-performing and most creative methods are invited for a second collaborative phase. The mandate of this phase is to work together to generate a tool with improved performance on the same problem.

The second stage of the DREAM challenges is essentially a selectively crowd-sourced collaboration and has a lot in common with other open-source initiatives to develop (rather than simply assess) collaborative software as part of a community event. BioHackathons, Neurosynth hackathons and other open-source programming meet-ups gather participants in a single location for a few days and encourage code sharing and cooperation to solve a bioinformatics problem. This approach can be a very productive way for researchers to share expertise and develop professional relationships, though it is less likely to work on problems that demand larger software engineering projects.

Open-source crowdsourcing can also follow a competitive model. For example, Sequence Squeeze was a race for the best genomic data compression algorithm and a cash prize. By submitting their code publicly as it evolved, participants received feedback on their performance on a leaderboard, spurring them to work harder to leapfrog over the current leader. A potential danger of revealing code is that many competitors will start basing their work on the leader's code at the expense of trying totally different directions, and feedback on performance can lead to overfitting on small data sets. An alternative could be to reveal code at one or two discrete points in the competition.

Although inspiring researchers to outpace their rivals for the top spot can yield positive results, competitions will benefit the community most if they take to heart Franklin D. Roosevelt's remark that "Competition has been shown to be useful up to a certain point and no further, but cooperation, which is the thing we must strive for today, begins where competition leaves off."